

Computational Protein Design Using AND/OR Branch-and-Bound Search

Yichao Zhou¹, Yuexin Wu¹, and Jianyang Zeng^{1,2,*}

¹ Institute for Interdisciplinary Information Sciences, Tsinghua University, P.R.C.

² MOE Key Laboratory of Bioinformatics, Tsinghua University, P.R.C.

Abstract. The computation of the global minimum energy conformation (GMEC) is an important and challenging topic in structure-based computational protein design. In this paper, we propose a new protein design algorithm based on the AND/OR branch-and-bound (AOBB) search, which is a variant of the traditional branch-and-bound search algorithm, to solve this combinatorial optimization problem. By integrating with a powerful heuristic function, AOBB is able to fully exploit the graph structure of the underlying residue interaction network of a backbone template to significantly accelerate the design process. Tests on real protein data show that our new protein design algorithm is able to solve many problems that were previously unsolvable by the traditional exact search algorithms, and for the problems that can be solved with traditional provable algorithms, our new method can provide a large speedup by several orders of magnitude while still guaranteeing to find the global minimum energy conformation (GMEC) solution.

Keywords: protein design, AND/OR branch-and-bound, global minimum energy conformation, residue interaction network, mini-bucket heuristic

1 Introduction

In a *structure-based computational protein design* (SCPD) problem, we aim to find a new amino acid sequence that accommodates certain structural requirements and thus can perform desired functions by replacing several residues from a wild-type protein template. The SCPD has exhibited promising applications in numerous biological engineering situations, such as enzyme synthesis [2], drug resistance prediction [9], drug design [13], and design of protein-protein interactions [28].

The aim of SCPD is to find the *global minimum energy conformation* (GMEC), that is, the global optimal solution of an amino acid sequence that minimizes a defined energy function. In practice, the rigid body assumption which anchors the backbone template is usually applied to reduce computational complexity. In addition, possible side-chain assignments for a residue are further discretized into several known conformations, called the *rotamer library*. It has been proved that SCPD is NP-hard [27] even with the two aforementioned prerequisites. A number of heuristic methods have been proposed to approximate the GMEC [30,20,24]. Unfortunately, these heuristic methods can be trapped

* Corresponding author: Jianyang Zeng. Email: zengjy321@tsinghua.edu.cn.

into local minima and may lead to poor quality of the final solution. On the other hand, several exact and provable search algorithms which guarantee to find the GMEC solution have been proposed, such as Dead-End Elimination (DEE) [6], A* search [21,22,7,35], tree decomposition [32], branch-and-bound (BnB) search [14,31,3], and BnB-based linear integer programming [1,18].

In our protein design scheme, a set of DEE criteria [12,10] is first applied to prune the infeasible rotamers that are provably not part of the GMEC solution. After that, the AND/OR branch-and-bound (AOBB) search [23] is used to traverse over the remaining conformational space and find the GMEC solution. In addition, we propose an elegant extension of this AND/OR branch-and-bound algorithm to compute the top k solutions within a user-defined energy cutoff from the GMEC. Our tests on real protein data show that our new protein design algorithm can address many design problems which cannot be solved exactly before, and for the problems that were solvable formerly, our new method can achieve a significant speedup by several orders of magnitude.

1.1 Related Work

The A* algorithm [21,17] uses a priority queue to store all the expanded states, which unfortunately may exceed the hardware memory limitation for large problems. AOBB, on the contrary, uses depth-first-search strategy that only requires linear space complexity with respect to the number of mutable residues.

The traditional BnB search algorithm [14] usually ignores the underlying topological information of the residue interaction network constructed based on the backbone template, while AOBB is designed to exploit this property.

Although the tree decomposition method [32] utilizes the residue interaction network, the table allocated by its dynamic programming routine may be too large to fit in memory. To fix this problem, AOBB adopts the mini-bucket heuristic to prune a large number of states to speeds up the search process.

2 Methods

2.1 Overview

Under the assumptions of rigid backbone structures and discrete side-chain conformations, the structure-based computational protein design (SCPD) can be formulated as a combinatorial optimization problem which aims to find the best rotamer sequence $r = (r_1, \dots, r_n)$ that minimizes following objective function:

$$E_T(r) = E_0 + \sum_{i=1}^n E_1(r_i) + \sum_{i=1}^n \sum_{j=i+1}^n E_2(r_i, r_j), \quad (1)$$

where n stands for the number of mutable residues, $E_T(r)$ represents the total energy of the system in which the rotamer assignment of the mutable residues

is r , E_0 represents the template energy (i.e., the sum of the backbone energy and the energy among non-mutable residues), $E_1(r_i)$ represents the self energy of rotamer r_i (i.e., the sum of intra-residue energy and the energy between r_i and non-mutable residues), and $E_2(r_i, r_j)$ is the pairwise energy between rotamers r_i and r_j .

2.2 AND/OR Branch-and-Bound Search

2.2.1 Branch-and-Bound Search Suppose we try to find the global minimum value of the energy function $E(r)$, in which $r \in R$ and R is the conformational space of the rotamers. The BnB algorithm executes two steps recursively. The first step is called *branching*, in which we split the conformational space R into two or more smaller spaces, i.e., R_1, R_2, \dots, R_m , where $R_1 \cup R_2 \cup \dots \cup R_m = R$. If we are able to find $\hat{r}_i = \arg \min_{r \in R_i} E(r)$ for all $i \in \{1, 2, \dots, m\}$, we can compute the minimum energy conformation in the conformational space R by identifying one of \hat{r}_i that has the lowest energy.

The second step of BnB is called *bounding*. Suppose the current lowest energy conformation is \bar{r}_i . For any sub-space R_j , if we can ensure that the lower bound of the energy of all conformations in R_j is greater than $E(\bar{r}_i)$, we can prune the whole sub-space R_j safely. The lower bound of the energy of the conformations in a given space usually can be computed based on some heuristic functions. The BnB algorithm performs the branching and bounding steps recursively until the current conformational space contains only one single conformation. A more detailed introduction to branch-and-bound search can be found in Appendix Section A1 [33].

2.2.2 Residue Interaction Network Traditional BnB algorithm can hardly exploit the underlying graph structure of the residue-residue interactions. In a real design problem, some mutable residues can be relatively distant and thus the pairwise energy terms in Equation (1) between these residues are usually negligible. Based on this observation, we can construct a *residue interaction network*, in which each node represents a residue, and two nodes are connected by an undirected edge if and only if the distance between them is less than a threshold. Fig. 1(a) gives an example of such a residue interaction network.

Consider a residue interaction network which contains two connected components (i.e., two clusters of mutable residues at two distant positions). Suppose each residue has at most p rotamers and the size of each connected component is q . Then the traditional BnB search needs to visit $O(p^{2q})$ nodes in the worst case. However, from the residue interaction network, we know that two connected components are independent, which means that altering the rotamers in one connected component does not affect the pairwise energy terms in the other connected component. So we can run the BnB search for each connected

component independently and then put the resulting minimum energy conformations together to form the GMEC solution, which only needs to visit $O(p^d)$ nodes in the worst case.

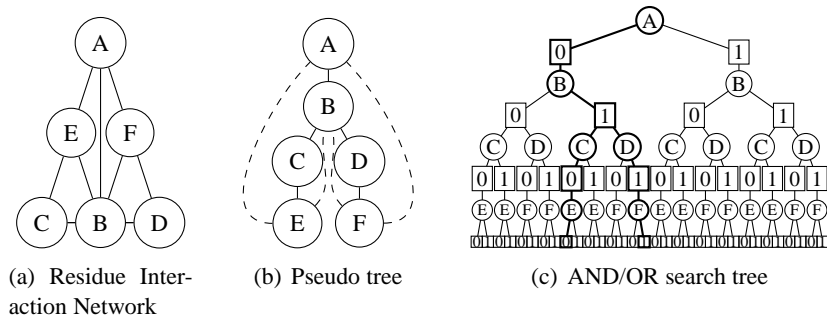


Fig. 1. An example of constructing an AND/OR search tree. (a) An example of a residue interaction network. (b) The corresponding pseudo-tree of the residue interaction network in (a), in which dashed lines are non-tree edges. (c) The full AND/OR search tree constructed from the pseudo-tree in (b), in which circle nodes represent OR nodes and rectangle nodes represent AND nodes. An example of a solution tree for the AND/OR search tree in (c) is marked in bold.

The independence requirement of connected components in a residue interaction network is too strict in practice. In fact, we can partition the whole network into several independent connected components after choosing particular rotamers in some residues. For example, after fixing the rotamers for residues *A* and *B* in the example shown in Fig. 1(a), we can obtain two independent components *CE* and *DF*. Then we can use the aforementioned method to reduce the size of search space and then search it using branch-and-bound algorithm. This is the major motivation of AND/OR branch-and-bound (AOBB) search [23].

2.2.3 AND/OR Search Space A *pseudo-tree* [8] of a graph (network) is a rooted spanning tree on that graph in which every non-tree edge in the graph is connected from a node to its offspring in the spanning tree. In other words, non-tree edges are not allowed to connect two nodes that are located in different branches of the spanning tree. Fig. 1(b) shows an example of a pseudo-tree constructed based on the residue interaction network in Fig. 1(a).

The pseudo-tree is a useful representation because for any node x in the tree, once all the side-chains of x and its ancestors are fixed, all the sub-trees rooted at the children of node x are independent. In other words, altering the rotamers for the sub-tree rooted at a child of x does not affect the total energy of the another sub-tree. Thus, the size of the search space for all sub-trees rooted at children of node x is proportional to the sum of the sizes of these sub-trees rather than the product of their sizes as in the traditional BnB algorithm. Therefore, AOBB often has a much smaller search space compared to the traditional BnB search.

The structure of an AOBB search tree is determined by its pseudo-tree. In order to represent the dependency between nodes, an AOBB search tree contains two types of nodes. The first type of nodes is called *OR nodes*, which splits the space into several parts that cover the original space by assigning a particular rotamer to a residue. The second type of nodes is called *AND nodes*, which decomposes the space into several smaller spaces where the computations of total energy of residues in different branches are independent to each other. The root of an AOBB search tree is an OR nodes and all the leaves are AND nodes. For each node in an AOBB search tree, its type is different from that of its parent. An example of an AOBB search tree is given in Fig. 1(c).

Unlike the traditional BnB search, in which a solution is represented by a single leaf node, in an AOBB search tree, a valid conformation is represented by a tree, called the *solution tree*. A solution tree shares the same root with the AOBB search tree. If an AND node is in the solution tree, all its OR children are also in the tree. If an OR node is in the solution tree, exact one of its AND children is in the tree. The tree with bold lines in Fig. 1(c) shows an example of a solution tree. In order to compute the best solution tree with the minimum energy when traversing the search space, we can maintain a *node value* $v(x)$ to store the total energy involving the residues in the sub-tree rooted at x . In an AOBB search tree, $v(x)$ can be computed as follows:

$$v(x) = \begin{cases} 0, & \text{if } x \text{ is a leaf node;} \\ \sum_{y \in \text{child}(x)} v(y), & \text{if } x \text{ is an internal AND node;} \\ \min_{y \in \text{child}(x)} e(y) + v(y), & \text{if } x \text{ is an internal OR node,} \end{cases} \quad (2)$$

where $\text{child}(x)$ stands for the set of children of node x and $e(y)$ is the sum of the self energy of the rotamer represented by y and the pairwise energy between the rotamer represented by y and other rotamers represented by the ancestors of y . Then the $v(\cdot)$ value of the root of the whole search tree is equal to the energy of the GMEC solution. The corresponding best solution tree can be constructed using a similar method.

Because AOBB uses the depth-first-search strategy, its space complexity is $O(n)$, where n is the number of mutable residues. The time complexity of AOBB in the worst case is $O(n * p^d)$, where p is the number of rotamers per residue and d is the depth of the pseudo-tree. A more detailed explanation about the AOBB search with pseudocode can be found in Appendix Section A2 [33].

2.2.4 Heuristic Function The choice of the heuristic function $h(x)$, which is a lower bound of $v(x)$, heavily affects the performance of the AOBB algorithm. A popular heuristic function used with AOBB is called *mini-bucket heuristic* [16], which is computed by the *mini-bucket elimination* algorithm [5]. The computation of mini-bucket heuristic can be accelerated through pre-computation, so

that $h(x)$ can be computed efficiently by looking up of pre-computed tables. The bound given by the mini-bucket heuristic can be further tighten by Max-Product Linear Programming [11] and Join Graph Linear Programming [15].

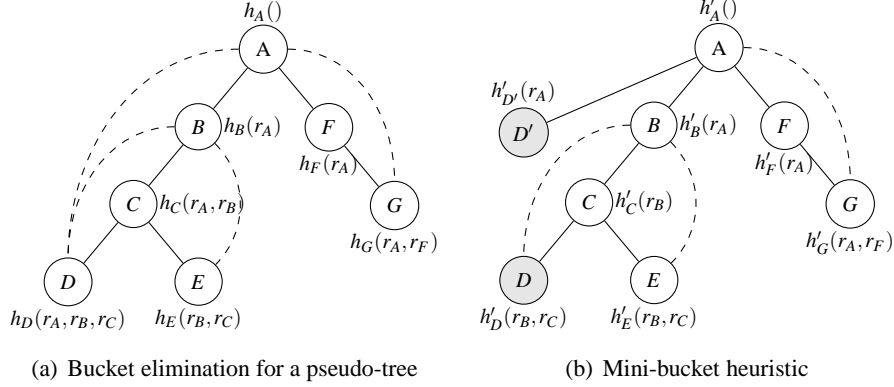


Fig. 2. An example of mini-bucket elimination. (a) The pseudo-tree of a graph along with the resulting energy tables computed by the bucket elimination algorithm. The dashed lines represents the non-tree edges in the original graph. (b) The tree generated by the mini-bucket elimination algorithm for the pseudo-tree in (a), in which the original energy table $h_D(r_A, r_B, r_C)$ is split into two smaller tables $h'_D(r_B, r_C)$ and $h'_{D'}(r_A)$.

The mini-bucket elimination is an approximation version of the *bucket elimination* algorithm [4], which is an another exact algorithm for solving the combinatorial problem with an underlying graph structure, such as protein design, based on a pseudo-tree. More specifically, the bucket elimination algorithm maintains an energy table $h_x(\cdot)$ for each tree node x , which stores the exact lower bound on the sum of energy involving the residues in the sub-tree rooted at x given the rotamer assignments of x 's ancestors. For instance, $h_D(r_A, r_B, r_C)$ in Fig. 2(a) stores the exact lower bound of node D given the rotamer assignments of its ancestors r_A , r_B , and r_C . These energy tables can be computed in a bottom-up manner. As an example, Fig. 2(a) shows the energy tables of the bucket elimination on a pseudo-tree of a residue interaction network, and we can compute $h_C(r_A, r_B) = \min_{r_C} (E(r_B, r_C) + h_D(r_A, r_B, r_C) + h_E(r_B, r_C))$, where $E(r_B, r_C)$ represents the pairwise energy term between rotamers r_B and r_C . The h value of the tree root, $h_A()$, in this example, is the total energy of the GMEC. The time complexity of bucket elimination is $O(n * \exp(w))$ [4], where n is the number of the nodes and w is the tree width [29] of the graph.

If the tree width of a graph is large, the energy tables may have high dimensions and thus can be too large to compute. The mini-bucket elimination is proposed to address this problem. In particular, it splits a node with a large en-

ergy table into multiple nodes with smaller energy tables, called *mini-buckets*, along with the pairwise energy term represented by the new added edges to decrease the dimension of its original energy table. We use $h'_x(\cdot)$ to represent the new energy table for each node x computed by the mini-bucket algorithm. Fig. 2(b) gives an example, in which $h_D(r_A, r_B, r_C)$ is split into two smaller tables $h'_D(r_B, r_C) = \min_{r_D} (E(r_D, r_B) + E(r_D, r_C))$ and $h'_{D'}(r_A) = \min_{r_D} E(r_D, r_A)$. Because now D and D' can be assigned with different rotamers, the new energy tables computed by the bucket elimination on the new graph is a lower bound of the original problem. Therefore, we can use the sum of $h'_x(\cdot)$ on all mini-buckets of a node as the heuristic function for AOBB.

2.3 Finding Sub-optimal Conformations

In practice, we often require the design algorithm to output the k best conformations within a given energy cutoff Δ [7]. In the BnB framework, this can be done easily by running the BnB search k times and remove the optimal conformations found in the preceding rounds from the search space. The task is more complicated to tackle in the AOBB because a conformation is represented by a solution tree rather than a tree node. Our solution consists of two parts:

1. In bounding steps, do not prune nodes in which the heuristic function values of the corresponding solution trees do not exceed the critical value by Δ .
2. Keep track of the k best solution trees and their $v(\cdot)$ values rather than only a single solution.

For the second part, we need to extend the procedure of computing $v(x)$, originally described in Equation (2). For each node x , we now store the k node values. Let $v_1(x)$ be the best node value, $v_2(x)$ be the second one, and so on. For each leaf node x , $v_1(x) = 0$ and $v_2(x) = v_3(x) = \dots = v_k(x) = \infty$. For each OR node x , we can compute $v_1(x) \leq v_2(x) \leq \dots \leq v_k(x)$ by merging $v_i(\cdot)$ values of x 's children using a sort routine and retaining the k smallest values.

The merge operation for AND nodes is very challenging. For each AND node x , let its children be y_1, y_2, \dots, y_t . Our task is to find k different sequences $(a_1, \dots, a_j, \dots, a_k)$, where $a_j = (a_{j1}, a_{j2}, \dots, a_{jt})$ and $a_{ji} \in \{1, 2, \dots, k\}$, so that $v_j(x) = \sum_{i=1}^t v_{a_{ji}}(y_i)$ and $v_1(x) \leq v_2(x) \leq \dots \leq v_k(x)$. A brute-force method for solving this problem requires $O(k^t)$ time complexity as it needs to enumerate all possible sequences for a_1, a_2, \dots, a_k , which is unacceptable because both k and t may be large in a real problem.

A simple example and the pseudocode of our merge operation for an AND node are shown in in Fig. 3. This algorithm uses a priority queue Q , which is a data structure that supports the operations of inserting a key/value pair (i.e., element) and extracting the element with the minimum value. We first define an index sequence $b = (b_1, \dots, b_t)$, where entry b_i represents the index of the chosen $v(\cdot)$ value in child y_i . Initially, $b = (1, 1, \dots, 1)$ is pushed to Q . In this

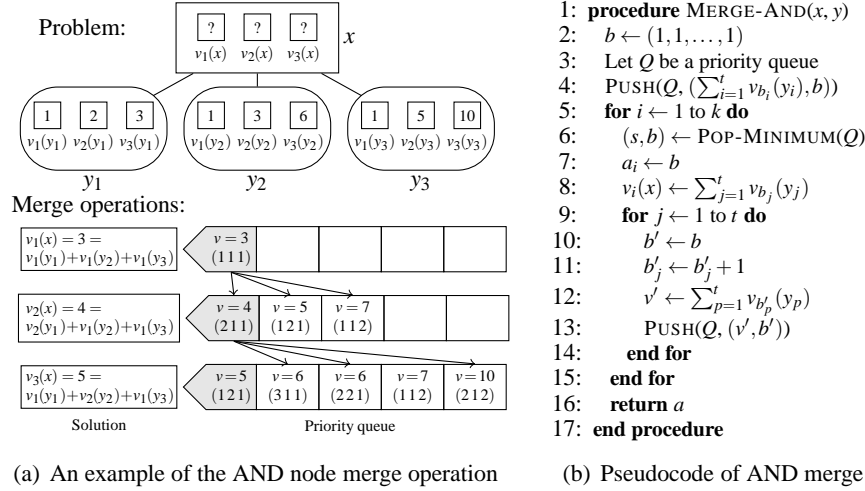


Fig. 3. The merge operation for AND nodes. (a) An example where the upper part describes the problem and the lower part shows how to solve this problem using a priority queue. The numbers in small squares show the corresponding $v(\cdot)$ values of individual tree nodes. The shaded boxes show the element with the smallest value in each priority queue. (b) The pseudocode of the merge operation for AND nodes.

problem, the value of an element is the sum of $v(\cdot)$ values of the AND nodes' children computed using the index sequence b as the key (Line 4). The initial index sequence $b = (1, 1, \dots, 1)$ corresponds to the first sequence a_1 because we choose the best v value for each child and thus we can get the best $v(\cdot)$ value for their parent. Each time we extract the element with the minimum value from Q as the next best sequence (Line 6). Then we push all the successors of the extracted sequence, computed by increasing only one index for each element in the sequence, into the priority queue (Lines 9 to 14). We repeat these steps until all the a_i values are generated. The time complexity of this process is $O(kt \log(kt))$. The proof of the correctness about our merge algorithm is provided in Appendix Section A3. [33]

3 Results

We conducted two computational experiments to evaluate the performance of our new AOBB-based protein design algorithm. In the first experiment, we compared our new AOBB-based algorithm with the traditional A*-based algorithm in a core redesign problem. To make a fair comparison, in this test we did not make any approximation in the energy matrix (i.e., the residue interaction network is fully connected) because the A*-based algorithm cannot benefit much from such approximation. In the second computational experiment, we performed the full protein design to examine the performance of our algorithm on a larger residue interaction network.

Our AOBB-based protein design algorithm was implemented based on the protein design package OSPREY [17] and the UAI branch of the AOBB search framework daoopt [25,26]. For comparison, we used the DEE/A* solver provided by the OSPREY package. In addition, we included the sequential A* solver with the improved computation of heuristic functions [34]. We used an Intel Xeon E5-1620 3.6GHz CPU in all evaluation tests.

3.1 Core Redesign

Core redesign can replace the amino acids in the core of a wild-type protein to increase its thermostability [19]. In this experiment, we tested all the 23 protein core redesign cases that failed to be solved in using the expanded rotamer library with the rigid DEE/A* in 4G memory from [10]. In addition, we picked another 5 design problems from [10] that were solvable within the given memory using the traditional DEE/A* algorithm. To make a fair comparison between A* and AOBB search algorithm, we did not remove any edge from the fully connected residue interaction network during the AOBB search in this test.

PDB	Space size	# of A* states	# of AOBB states	OSPREY time	cOSPREY time	AOBB time
1TUK	1.73e+19	OOM	188,042	OOM	OOM	723
1ZZK	3.44e+15	OOM	255	OOM	OOM	< 1
2BWF	5.54e+22	OOM	517,258,245	OOM	OOM	1,467,951
3FIL	2.62e+21	OOM	3	OOM	OOM	< 1
2RH2	1.29e+22	OOM	OOT	OOM	OOM	OOT
1IQZ	7.11e+17	18,337,117	90,195	1,824,235	40,217	117
2COV	1.14e+10	43,306	3	317	21	1
3FGV	6.44e+12	3,073,965	3	59,589	5,091	< 1
3DNJ	5.11e+12	569,597	4,984	7,469	570	3
2FHZ	1.83e+18	14,732,913	3,972	3,475,716	70,783	13

Table 1. The comparison between A*-based and AOBB-based algorithms on core redesign

Table 1 summarizes the comparison results between A*-based and our AOBB-based algorithms, in which “OOM” and “OOT” represents “out of memory” and “out of time”, respectively. The full comparison results can be found in Appendix Table A1 [33]. The memory was limited to 4G, which was the same as that in [10], and the running time was limited to 8 hours. The first five rows show the five cases (among 23 cases) in [10] which were formerly unsolvable by the original A* algorithm. The column labeled with “Space size” shows the size of the conformational space after DEE pruning. The columns labeled with “OSPREY time” and “cOSPREY time” show running time of the A* solvers from OSPREY and [34], respectively. The running time was measured in millisecond and did not include the initialization steps of each algorithm. The initialization time of AOBB was relatively stable for all cases and typically took 90s to compute the mini-bucket heuristic tables and an initial bound for AOBB search.

As shown in Table 1 and Table A1, the AOBB algorithm can successfully find the GMEC solutions for 21 out of the 23 problems from [10]’s data, which

were formerly unsolvable by the original A* algorithm in 4G memory. Also, we find that the number of states expanded in the AOBB search was much less than that in the traditional A* search. Accordingly, for those cases solvable by both algorithms, the AOBB search consumed less time than the traditional A* search. Probably this improvement was due to the fact that the mini-bucket heuristic with MPLP and JGLP is tighter than the heuristic function used in OSPREY.

3.2 Full Protein Design

In the second computational experiment, we ran the full protein design to evaluate the performance of our AOBB-based protein design algorithm. In the full protein design problem, all residues of a protein are mutable, which leads to a much larger conformational space. For each residue, we picked 1-4 the most similar amino acids, according to the BLOSUM62 matrix, as the mutation candidates. For each pair of residues A and B , we added an edge (A, B) to the residue interaction network if and only if for all rotamer assignments r_A and r_B , $(\max_{r_A, r_B} E(r_A, r_B) - \min_{r_A, r_B} E(r_A, r_B)) > \lambda$, where threshold parameter λ was used to trade the precision of the energy with the easiness of the problem. We used $\lambda = 0.04$ for all the test cases.

PDB	Space size	# of residues	# of edges	Tree depth	# of AOBB states	AOBB time
1I27	6.69e+45	69	968	40	3,149	11
1M1Q	2.33e+19	71	390	17	3	< 1
1T8K	2.83e+43	75	1031	42	3	< 1
1XMK	2.66e+48	74	1108	40	864	2
3G36	4.28e+20	47	396	22	159	< 1
3JTZ	1.96e+45	71	961	44	4,354,110	17,965

Table 2. The test results on the full protein design problem

Table 2 shows the test results of this computational experiment. The running time was measured in millisecond. Here we did not list the results of traditional A*-based algorithm because we found that A*-based algorithms were unable to find the GMEC solutions for all these test cases within 4G memory. The AOBB-based search algorithm can found the GMEC solutions for all the test cases. This demonstrates the power of the AOBB search algorithm with the state-of-the-art heuristic function, which can effectively address full protein design problems.

4 Conclusion and Future Work

In this paper, we developed a new protein design algorithm based on the new branch-and-bound search technique (i.e., AOBB) to find the global minimum energy conformation, which speeds up the search process by several orders of magnitude than the traditional provable algorithms. The AOBB-based algorithm accelerates the search process based on an advanced heuristic function and fully exploits the topology of the residue interaction network while it only has linear

memory consumption. The algorithm can also output suboptimal solutions by employing an elegant modification of the original search algorithm.

Currently, our algorithm is only implemented on a single machine. It is possible to further accelerate the design process by parallelizing AOBB search on a GPU processor or a CPU cluster on a supercomputer, which will enable us to deal with the protein design problems with larger conformational space.

Acknowledgement

We thanks Dr. Lars Otten and Prof. Alex Ihler for their support in providing their code of the daoopt AOBB solver.

Funding: This work was supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003 and 61472205, and China’s Youth 1000-Talent Program.

References

1. E. Althaus, O. Kohlbacher, H.-P. Lenhof, and P. Müller. A combinatorial approach to protein docking with flexible side chains. *Journal of Computational Biology*, 9(4):597–612, 2002.
2. C.-Y. Chen, I. Georgiev, A. C. Anderson, and B. R. Donald. Computational structure-based redesign of enzyme activity. *Proceedings of the National Academy of Sciences*, 106(10):3764–3769, 2009.
3. J. D. D. Allouche, G. K. S. de Givry, I. A. T. Schiex, S. T. S. Barbe, and B. O. S. Prestwich. Computational protein design as an optimization problem. *Artificial Intelligence*, 212:59–79, 2014.
4. R. Dechter. Bucket elimination: A unifying framework for probabilistic inference. In *Learning in Graphical Models*, pages 75–104. Springer, 1998.
5. R. Dechter and I. Rish. Mini-buckets: A general scheme for bounded inference. *Journal of the ACM (JACM)*, 50(2):107–153, 2003.
6. J. Desmet, M. D. Maeyer, B. Hazes, and I. Lasters. The dead-end elimination theorem and its use in protein side-chain positioning. *Nature*, 356(6369):539–542, 1992.
7. B. R. Donald. *Algorithms in structural molecular biology*. The MIT Press, 2011.
8. E. C. Freuder and M. J. Quinn. Taking advantage of stable sets of variables in constraint satisfaction problems. In *International Joint Conference on Artificial Intelligence*, volume 85, pages 1076–1078, 1985.
9. K. M. Frey, I. Georgiev, B. R. Donald, and A. C. Anderson. Predicting resistance mutations using protein design algorithms. *Proceedings of the National Academy of Sciences*, 107(31):13707–13712, 2010.
10. P. Gainza, K. E. Roberts, and B. R. Donald. Protein design using continuous rotamers. *PLoS Computational Biology*, 8(1):e1002335, 2012.
11. A. Globerson and T. S. Jaakkola. Fixing max-product: Convergent message passing algorithms for MAP LP-relaxations. In *Advances in neural information processing systems*, pages 553–560, 2008.
12. R. F. Goldstein. Efficient rotamer elimination applied to protein side-chains and related spin glasses. *Biophysical Journal*, 66(5):1335–1340, 1994.
13. M. J. Grczynski, J. Grembecka, Y. Zhou, Y. Kong, L. Roudaia, M. G. Douvas, M. Newman, I. Bielnicka, G. Baber, T. Corpora, et al. Allosteric inhibition of the protein-protein interaction between the leukemia-associated proteins Runx1 and CBF β . *Chemistry & biology*, 14(10):1186–1197, 2007.
14. E.-J. Hong and T. Lozano-Pérez. Protein side-chain placement through MAP estimation and problem-size reduction. In *Algorithms in Bioinformatics*, pages 219–230. Springer, 2006.

15. A. T. Ihler, N. Flerova, R. Dechter, and L. Otten. Join-graph based cost-shifting schemes. *arXiv preprint arXiv:1210.4878*, 2012.
16. K. Kask and R. Dechter. A general scheme for automatic generation of search heuristics from specification dependencies. *Artificial Intelligence*, 129(1):91–131, 2001.
17. D. A. Keedy, C.-Y. Chen, F. Rezam, and A. C. Andersonl. OSPREY: Protein design with ensembles, flexibility, and provable algorithms. *Methods in Protein Design*, page 87, 2013.
18. C. L. Kingsford, B. Chazelle, and M. Singh. Solving and analyzing side-chain positioning problems using linear and integer programming. *Bioinformatics*, 21(7):1028–1039, 2005.
19. A. Korkegian, M. E. Black, D. Baker, and B. L. Stoddard. Computational thermostabilization of an enzyme. *Science*, 308(5723):857–860, 2005.
20. B. Kuhlman and D. Baker. Native protein sequences are close to optimal for their structures. *Proceedings of the National Academy of Sciences*, 97(19):10383–10388, 2000.
21. A. R. Leach, A. P. Lemon, et al. Exploring the conformational space of protein side chains using dead-end elimination and the A* algorithm. *Proteins Structure Function and Genetics*, 33(2):227–239, 1998.
22. S. M. Lippow and B. Tidor. Progress in computational protein design. *Current Opinion in Biotechnology*, 18(4):305–311, 2007.
23. R. Marinescu and R. Dechter. AND/OR branch-and-bound search for combinatorial optimization in graphical models. *Artificial Intelligence*, 173(16):1457–1491, 2009.
24. J. S. Marvin and H. W. Hellenga. Conversion of a maltose receptor into a zinc biosensor by computational design. *Proceedings of the National Academy of Sciences*, 98(9):4955–4960, 2001.
25. L. Otten and R. Dechter. Anytime and/or depth-first search for combinatorial optimization. *AI Communications*, 25(3):211–227, 2012.
26. L. Otten, A. Ihler, K. Kask, and R. Dechter. Winning the PASCAL 2011 MAP challenge with enhanced AND/OR branch-and-bound. In *DISCML*, 2012.
27. N. A. Pierce and E. Winfree. Protein design is NP-hard. *Protein Engineering*, 15(10):779–782, 2002.
28. K. E. Roberts, P. R. Cushing, P. Boisguerin, D. R. Madden, and B. R. Donald. Computational design of a PDZ domain peptide inhibitor that rescues CFTR activity. *PLoS Computational Biology*, 8(4):e1002477, 2012.
29. N. Robertson and P. D. Seymour. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986.
30. A. G. Street and S. L. Mayo. Computational protein design. *Structure*, 7(5):R105–R109, 1999.
31. S. Traoré, D. Allouche, I. André, S. de Givry, G. Katsirelos, T. Schiex, and S. Barbe. A new framework for computational protein design through cost function network optimization. *Bioinformatics*, 29(17):2129–2136, 2013.
32. J. Xu and B. Berger. Fast and accurate algorithms for protein side-chain packing. *Journal of the ACM (JACM)*, 53(4):533–557, 2006.
33. Y. Zhou, Y. Wu, and J. Zeng. Appendix of “computational protein design using AND/OR branch-and-bound search”. Available at <http://iiis.tsinghua.edu.cn/~compbio/papers/recomb15A0BBapx.pdf>, 2015.
34. Y. Zhou, W. Xu, B. R. Donald, and J. Zeng. An efficient parallel algorithm for accelerating computational protein design. *Bioinformatics*, 30(12):i255–i263, 2014.
35. Y. Zhou and J. Zeng. Massively parallel A* search on a GPU. In *Proceedings of the National Conference on Artificial Intelligence*, 2015.